Improving Ambulance Dispatching with Machine Learning and Simulation

Citation for published version (APA):

Theeuwes, N., van Houtum, G. J., & Zhang, Y. (2021). Improving Ambulance Dispatching with Machine Learning and Simulation. In Y. Dong, N. Kourtellis, B. Hammer, & J. A. Lozano (Eds.), *Machine Learning and Knowledge Discovery in Databases: Applied Data Science Track : European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part IV* (pp. 302-318). (Lecture Notes in Computer Science (LNCS); Vol. 12978), (Lecture Notes in Artificial Intelligence (LNAI); Vol. 12978). Springer. https://doi.org/10.1007/978-3-030-86514-6_19

DOI:

[10.1007/978-3-030-86514-6_19](https://doi.org/10.1007/978-3-030-86514-6_19)

Document status and date:

Published: 01/01/2021

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Improving ambulance dispatching with machine learning and simulation

Nikki Theeuwes, Geert-Jan van Houtum^[0000–0003–4225–5434], and
Yingqian Zhang (✉)^[0000–0002–5073–0787]

Eindhoven University of Technology, 5612 AZ Eindhoven, Netherlands
`nikkithewes@gmail.com`, `{G.J.v.Houtum,yqzhang}@tue.nl`

Abstract. As an industry where performance improvements can save lives, but resources are often scarce, emergency medical services (EMS) providers continuously look for ways to deploy available resources more efficiently. In this paper, we report a case study executed at a Dutch EMS region to improve ambulance dispatching. We first capture the way in which dispatch human agents currently make decisions on which ambulance to dispatch to a request. We build a decision tree based on historical data to learn human agents’ dispatch decisions. Then, insights from the fitted decision tree are used to enrich the commonly assumed closest-idle dispatch policy. Subsequently, we use the captured dispatch policy as input to a discrete event simulation to investigate two enhancements to current practices and evaluate their performance relative to the current policy. Our results show that complementing the current dispatch policy with redispatching and reevaluation policies yields an improvement of the on-time performance of highly urgent ambulance requests of 0.77 percentage points. The performance gain is significant, which is equivalent to adding additional seven weekly ambulance shifts.

Keywords: Ambulance dispatching · Machine learning · Decision trees · Discrete event simulation · Logistics

1 Introduction

Emergency medical services (EMS) providers continuously look for ways to deploy limited available resources more efficiently. In the Netherlands, the fraction of highly urgent ambulance requests (A1 requests) with a response time of fewer than 15 minutes has been consistently below the national target of 95% throughout the past years, with a performance of 92.4% in 2017. Advances in ambulance logistics will contribute to the provision of sufficient emergency medical care, given the available resources.

The operational problems in EMS literature include both ambulance dispatching and relocation in order to maximize the fraction of ambulance requests with a *response time* below a certain threshold time, or the *on-time performance*. Response time is defined as the time between the moment an ambulance request arrives at a dispatch center and the moment the ambulance arrives at the request

location. It is predominantly assumed that ambulances are dispatched according to a ‘closest-idle’ policy (e.g. [10, 13]). Alternative dispatch policies are often modifications of this policy (e.g. [7, 8]). However, not only does this policy neglect practical considerations (e.g. the distinction between urgency levels or shift ends), it is also known to be suboptimal when maximizing the on-time performance [4]. Therefore, it can be expected that in practice dispatch agents tend to deviate from this commonly assumed dispatch policy, jeopardizing the relevance of alternative policies developed with the closest-idle policy at its foundation.

Knowledge obtained by dispatch agents in practice can be very useful in the development of improved dispatching policies [1]. In this paper, we formally capture this knowledge, or expertise, in the form of the current dispatch policy for the Dutch EMS region of Brabant Zuid-Oost (BZO). Capturing the current dispatch policy has three main benefits: (1) *Creating transparency*: Insights can be deduced that can help create awareness among dispatch agents, which might improve consistency and fairness of the process. (2) *Improving process*: Insights from the captured dispatch policy also give rise to opportunities for improvement. The captured dispatch policy provides a basis to improve upon by extending it with a number of additional or adapted decision rules. Contrary to developing an improved dispatch policy from scratch, our approach complements, rather than replaces, current dispatch practices. This ensures both the incorporation of practical considerations in the resulting policy and that it is in line with the way in which dispatch agents currently work, which are expected to foster adoption in practice; (3) *Evaluating fairly*: the captured dispatch policy can be used as input to a simulation of an EMS region. Such a simulation can also be used to fairly evaluate potential improvements of the dispatch process by comparing its performance to that of the current dispatch policy. The use of a benchmark that resembles current practices allows for more accurate conclusions regarding the potential of the evaluated alternative policy in practice.

Decision makers are often not completely aware of the reasoning behind their expert judgments, making it hard for them to verbally express their decision process [6]. However, mental decision models can be formally approximated through machine learning models. We select *decision tree induction* to capture the current dispatch policy of the BZO region, since this method results in a policy representation that is both transparent and interpretable. In the BZO region the fraction of A1 requests has been consistently below the nationally-set target of 95% (i.e., 91.7% in 2017), while the fraction of moderately urgent (A2) requests with a response time of less than 30 minutes has consistently exceeded its target of 95%. Therefore, the on-time performance for A1 requests is generally regarded as the main performance measure in EMS management.

Our work contributes to the field of EMS management as well as to that of applying machine learning to capture expert decisions:

- We are the first to formally capture current ambulance dispatch practices using machine learning. We apply decision tree induction to obtain a transparent representation of the current dispatch decision process in the BZO region (see §4.1).

- We apply a unique post-processing phase which combines knowledge from both practice and literature with the learned decision tree to further improve the quality of the learned model in terms of accuracy and conciseness. The resulting model enriches the commonly assumed closest-idle dispatch policy through the use of penalty values that reflect the risk associated with certain ambulance characteristics (§4.2).
- We illustrate an application of the captured current dispatch policy by proposing two enhancements to it and evaluating these in a simulation using the captured policy as a practically relevant benchmark (§5).

Before making these contributions in §4 and §5, we discuss related literature in §2 and the collection of data in §3. We conclude this paper in §6.

2 Related work

The existing studies in EMS management generally evaluate the proposed dispatch policies through a simulation in which many simplifying modelling choices and assumptions are made. For example, Lee [7] simulated a hypothetical square grid of 25 vertices with a fixed driving time for all edges. He did not distinguish between urgency levels and assumed a general distribution for transfer times and a static number of ambulances. Jagtenberg et al. [4] simulated the actual EMS region of Utrecht, but assumed a static relocation policy, static request arrivals, and static ambulance capacity, treatment and transfer times.

The existing, limited number of studies applying machine learning to model expert decisions generally seems to have the captured expert knowledge as the ultimate goal of their efforts, mostly to automate decision making. Maghrebi et al. [9] conducted a feasibility study of automating the process of determining the order of concrete deliveries. They employ machine learning to match expert decisions with the objective of decreasing dependency on human resources. Lafond et al. [5] compare three machine learning techniques in capturing human classification behavior using a simulated naval air defense task. However, capturing expert decisions with the objective to support future decisions implicitly assumes that the captured expert knowledge is optimal, or at least neglects the fact that insight into current practices provides a good opportunity for the identification and evaluation of improvement of the decision making process. In Lafond et al. [6], a learning technique is applied to functionally mirror expert mental models. Their objective is to improve decision quality by recognizing when a decision maker is deviating from his usual decision patterns, since this might indicate probable errors. It still assumes the captured policy to be the correct, or desired one, which was a limitation as acknowledged by the authors. Donnot et al. [3] first apply a deep neural network to historic decision data to mimic human decisions in the prevention of violating power flow limits in a power plant, and then use simple simulation to evaluate the effect of each action proposed by the captured decision model before suggesting it to the decision maker. While this approach does not actually improve on the captured decisions, it does distin-

guish between bad and good decisions and only uses the good ones to support future decision making.

To the best of our knowledge, there are no studies which have captured expert decisions with the objective to use the resulting policy as a basis to improve upon or as input for fair evaluation of alternative policies. Moreover, most of the studies did not derive decisions from real data, but rather generated this data by presenting experts with an artificial (simulated) task. In comparison, we expect decisions derived from historic data resemble actual decisions more closely. In addition, in capturing ambulance dispatch decisions, we apply a post-processing phase which combines knowledge from both the domain and literature with the learned model to further improve the quality the resulting model.

3 The data set: historic dispatch decisions

We approached the induction of the current dispatch policy as a classification problem. We gathered data on historic dispatch decisions made by BZO’s dispatch agents. The data set has been compiled such that it reflects all information available to the agent at the decision moment, which might have affected the decision. We have structured the decision to be captured around the *dispatch proposal*. In the Netherlands, upon being presented with an ambulance request, a dispatch agent uses the national dispatch system to generate such a dispatch proposal. A dispatch proposal is an ordered list of all ambulances available for dispatch to the concerned request, based on an increasing driving time to the request location. By structuring the decision to be captured around such a dispatch proposal, we have implicitly assumed that, for any dispatch decision to be made, a dispatch proposal is generated and one of the ambulances in the proposal is dispatched. The set of ambulances available for dispatch depends on the request’s urgency. Regardless of the request’s urgency this set includes all idle ambulances, i.e. those driving to, or waiting at, a station. Besides idle ambulances, this set includes ambulances which have already been dispatched to a less urgent request, but did not arrive at that request’s location yet, and ambulances that have arrived at a hospital and are busy transferring a patient. While these ambulances are not idle (yet), they might be redispached or requested to accelerate the transfer process respectively. Lastly, since dispatch agents have the possibility to request assistance from neighbouring EMS regions, these ambulances are also included in the dispatch proposal. Summarizing, the objective of our formalization effort was to determine which ambulance is dispatched to a request, given the corresponding dispatch proposal, and why a dispatch agent might decide to deviate from dispatching the closest-idle ambulance. This implies that the *class* of each *instance* is the rank of the ambulance that was actually dispatched in the corresponding dispatch proposal.

3.1 Feature engineering

Upon making a dispatch decision, a dispatch agent has multiple screens at his/her disposal which show information regarding the concerned ambulance re-

Table 1. Features with $i \in \{1, 2, 3, 4, 5\}$ being dispatch proposal options

No.	Feature	Symbol	Data type
0	Rank of i in dispatch proposal (class)	C	Nominal: $\{1, 2, 3, 4, 5+\}$
1	Urgency	U	Ordinal: $\{A1, A2\}$
2	Passed time	P	Numeric (minutes)
3-7	Driving time of i	D_i	Numeric (minutes)
8-12	Status of i	S_i	Nominal: $\{1, 2, 3, 6\}$
13-17	Idle status indicator of i	SI_i	Binary
18-22	Status time of i	ST_i	Numeric (minutes)
23-27	Remaining shift time of i	RS_i	Numeric (minutes)
28-32	Own ambulance indicator of i	$Rown_i$	Binary
33-37	Region BZO & BNO indicator of i	Rbo_i	Binary
38	Number of idle ambulances	I	Numeric
39	Single coverage	Cov	Numeric (%)
40-44	Percentual coverage reduction of i	PCR_i	Numeric (%)
45-49	Absolute coverage reduction of i	ACR_i	Numeric (%)
50-53	Driving time diff. i and $i + 1$	ΔD_i	Numeric (Min.)
54-57	Perc. coverage reduction diff. i and $i + 1$	ΔPCR_i	Numeric (%)
58-61	Abs. coverage reduction diff. i and $i + 1$	ΔACR_i	Numeric (%)
62-66	Expected response time of i	E_i	Numeric (Min.)

quest, the ambulance options included in the generated dispatch proposal, and a map of the region displaying all on-duty ambulance locations and statuses. Since dispatch agents are dedicated to making dispatch decisions, which happens under time pressure, we assume that all information presented to a dispatch agent is considered to be relevant to the dispatch decision. We transformed such domain knowledge by the process of feature engineering. Data was obtained for September and October 2018 from GMS. Only data on ambulance requests (dispatches) within the BZO region were used. This led to a total of 4506 instances to fit BZO's current dispatch policy on. Table 1 shows the features for which values were obtained from the available data for each of the instances. For a more detailed description of data collection and (pre-)processing, we refer to [12].

The first two features relate to the ambulance request which requires a dispatch decision. The urgency (U) of this request is relevant since it determines the response time target. 95% of highly urgent ($A1$) requests should have a response time less than fifteen minutes, while 95% of moderately urgent ($A2$) requests should have a response time less than thirty minutes. Since the response time of a request starts at the moment the corresponding call arrives at the dispatch center, the time that has passed since call arrival (P) is also relevant.

Features three to thirty-seven concern pieces of information listed for each of the ambulance options in the generated dispatch proposal, with i referring to the i th option in a dispatch proposal, $i \in \{1, 2, 3, 4, 5\}$. Note that for each instance only features referring to properties of the first five options in the concerned proposal are included. This choice was made since the class distribution in our instance set is particularly unbalanced, with the higher ranked dispatch options being represented more strongly. Recall that ambulances in a dispatch proposal are ordered based on their driving time to the concerned incident and the main performance measure depends strongly on this driving time, which leads to a natural preference for higher ranked options. To ensure a sufficient number of samples of each class to be available, classes five and up were combined to form

one class. Furthermore, we were especially interested in an agent’s reasons for deviating from sending the closest idle ambulance, which were expected to become apparent by distinguishing between the first few options of a dispatch proposal. The resulting class distribution is: 67% (class 1), 20% (class 2), 7% (class 3), 3% (class 4), 3% (class 5). Since the importance of the classes is ordered, it is more important that the higher ranked classes are predicted correctly. Hence we do not balance the dataset but let the decision tree algorithm favour the more important classes during learning.

The status of each dispatch option i (S_i) may either be idle (driving towards or waiting at a station) or busy but available for dispatch (on its way to a less urgent request or transferring a patient at a hospital). The idle status indicator of dispatch option i (SI_i) indicates whether S_i is idle. Furthermore, the status time of option i (ST_i) is equal to the time since the status of each dispatch option last changed, while the time until the end of each dispatch option’s eight hour shift, which may be negative in case of overtime, is reflected by feature RS_i . The dispatch proposal shows for each option to which region it belongs, and thus by which region it is controlled. We captured this information in binary features $Rown_i$ and Rbo_i , where the first reflects whether option i belongs to the own region (BZO), and the second indicates whether option i belongs to either the own region or the adjacent BNO region, where dispatch agents operate from the same dispatch center as BZO’s dispatch agent.

Features thirty-eight through forty-nine reflect the information that the dispatch agent might deduct from the map of the region displaying all on-duty ambulance locations and statuses. The number of idle ambulances (I) and the single coverage (Cov) reflect the extent to which the region is prepared for future requests. Based on discussions with BZO’s dispatch agents, I includes both idle ambulances and ambulances that are busy transferring a patient at a hospital, since these are expected to become idle in the very near future and may even be requested to accelerate the transfer process if necessary. The single coverage feature refers to the fraction of the BZO region (in terms of 4-digit postal code areas) that can be reached within a response time of fifteen minutes by at least one ambulance [2]. Additionally, we introduced two features that relate to the reduction in preparedness, i.e. single coverage, of the region that would be caused by dispatching option i . ACR_i does so in absolute terms, while PCR_i relates the coverage reduction to the current single coverage (Cov).

A dispatch agent might infer relevant information based on the relation between feature values. Features fifty through sixty-six were constructed by performing logical operations on our initial list of features and selecting meaningful ones. These additional features include the difference between subsequent dispatch options in driving time (ΔD_i), percentual and absolute coverage reduction (ΔPCR_i and ΔACR_i), and the expected response time of each dispatch option i (E_i). Here, the expected response time of option i is made up of the time that passed since arrival of the call (P), its driving time to the request’s location (D_i), and one minute that is expected to be required for making the dispatch decision and for an ambulance to start driving after being dispatched.

4 Capturing the dispatch policy with a decision tree

We use a decision tree to learn the current dispatch process, due to its transparent nature. Its interpretability allows us to gain insight into the current dispatch routine, which can be leveraged both as a basis to improve the current dispatch process and as a benchmark in the evaluation of potential improvements.

We split the data into a training set (70% of instances) and a test set (remaining 30% of instances). We use the implementation of CART (Classification and Regression Trees) in scikit-learn [11]. We tune the parameters, i.e. feature selection method, maximum tree depth, and the minimum number of instances at a leaf node, by applying stratified 10-fold cross-validation on the training set. Then the final decision tree has been trained on the complete training set. Subsequently, the resulting decision tree has been evaluated using the test set.

Since our objective of capturing the current dispatch process is to identify which ambulance is actually dispatched, the larger sized classes are of greater interest than the smaller ones. By definition, this relative interest in correctly predicting each class is reflected in the class distribution. Hence, we do not balance the training set but let the algorithm favour the more important classes. In addition, we choose the weighted F1-score, where the F1-score of each class is weighted by its sample size, as the main performance measure.

4.1 Performance analysis of the learned decision tree and policy

Additionally, we define the *Weighted Mean Error* performance measure. For the problem at hand, if the actual dispatch decision was to dispatch the first option, predicting dispatch of the third option is actually more wrong than predicting dispatch of the second option. Therefore, we defined the following additional performance measure:

$$WME = \frac{\sum_{d=0}^{k-1} d \sum_{i,j \in \{1,2,\dots,k\}: |i-j|=d} m_{i,j}}{\sum_{i,j \in \{1,2,\dots,k\}} m_{i,j}},$$

where k equals the number of possible classes, in our case $k = 5$, and the $m_{i,j}$ are cells in the confusion matrix, where rows and columns are indicated by i and j respectively. Naturally, while we strive towards a dispatch prediction model with a weighted F1-score that is as high as possible, we prefer the mean distance to the actual class to be as low as possible.

To place the performance of the resulting decision tree into perspective, its performance has been compared to the dispatch policy that is commonly assumed in literature, the *closest-idle policy*. Notice that in literature this policy generally does not include the additional dispatch options that are available to BZO's dispatch agents, namely ambulances that are not completely idle but nevertheless available to (certain) incidents and external ambulances that belong to other regions. Therefore, we have defined two dispatch policies to which the performance of our fitted dispatch policy have been compared: (1) *The limited closest-idle policy*: corresponding to the policy that is commonly assumed

in literature, i.e. dispatching the highest ranked ambulance in the dispatch proposal that is completely idle (on the road or at station) and belongs to the own region. (2) *The extended closest-idle policy*: corresponding to the commonly assumed policy but adapted to include the additional available dispatch options, i.e. always dispatching option one in the dispatch proposal.

Figure 1 depicts the learned decision tree. Figures 2a, b, and c show the confusion matrices and performance measures for the learned dispatch policy, the limited closest-idle policy and the extended closest-idle policy respectively. Figure 2 shows that the learned dispatch model outperforms both interpretations of the closest-idle policy, in terms of the weighted F1-score, as well as the weighted mean error. However, while the difference in performance between the learned model (a) and the extended closest-idle policy (c) is quite significant, the improvement in predictive performance of the learned model (a) relative to the basic, limited closest-idle policy (b) is less apparent. This observation leads us to believe that BZO’s dispatch agents generally make limited use of the additional dispatch options available to them.

This insight is confirmed by studying the learned decision tree, depicted in Figure 1a, in more detail. There are several clear ‘decision paths’, which have been highlighted in Figure 1b. These highlighted decision paths indicate the dominant dispatch decision. Note that some of these paths, and the insights derived from them, can be regarded as more important than others due to the larger number of samples following that path. The weight of each path indicates the number of samples following that path.

The main reasons that might lead a dispatch agent to deviate from dispatching the highest ranking dispatch option (i.e. option 1) quickly become clear from the splits on the most dominant path (leading to [A]). These main reasons include this highest ranking ambulance:

- **Not being immediately available for dispatch**: due to its status. For example, the ambulance is transferring a patient at a hospital, meaning that it might require some time to be relieved from its current request and redispached to the new request.
- **Not belonging to the own region**: meaning that the concerned dispatch center needs to be requested, which takes time, and the dispatch request might be denied.
- **Nearing the end of its shift**: causing a risk of overtime if it is dispatched.

The first two of these reasons confirm that dispatch agents make limited use of the additional dispatch options available to them. Possibly, this is the case because these issues add a potential delay to the indicated driving time. Such a potential delay adds a degree of uncertainty to the ambulance’s expected driving time, which gives the dispatch agent good reasons to deviate from this option. Naturally, the potential delay is only relevant if the difference between the driving time of that option and the subsequent option is less than this expected delay. This is reflected by the node at the top of node group [A] in Figure 1b, as well as at several other nodes in the tree.

Fig. 1. Visualization of learned dispatch decision model with colors indicating the class distribution of the instances reaching each node: (a) complete model and (b) model including highlighted decision paths indicating its dominant dispatch decision

		Predicted class				
		1	2	3	4	5
Actual class	1	873	22	12	7	0
	2	85	177	12	1	0
	3	29	31	26	5	0
	4	13	11	8	4	2
	5+	12	7	5	2	8
Weighted F1-score: 78.6%						
Weighted mean error: 0.30						
(a)						

		Predicted class				
		1	2	3	4	5
Actual class	1	837	22	16	7	32
	2	98	160	4	3	10
	3	25	25	36	2	3
	4	13	12	1	10	2
	5+	9	4	5	4	12
Weighted F1-score: 77.5%						
Weighted mean error: 0.40						
(b)						

		Predicted class				
		1	2	3	4	5
Actual class	1	914	0	0	0	0
	2	275	0	0	0	0
	3	91	0	0	0	0
	4	38	0	0	0	0
	5+	34	0	0	0	0
Weighted F1-score: 54.5%						
Weighted mean error: 0.52						
(c)						

Fig. 2. Confusion matrices and performance measures for (a) the learned dispatch policy, (b) the limited closest-idle policy, and (c) the extended closest-idle policy

It can be deduced that, if there are enough reasons to deviate from the highest ranking ambulance option, the subsequent option is considered. However, the same reasons to deviate seem to hold for this option, e.g. see the path in Figure 1b leading to node [B], where option 3 is considered due to the status of option 2, and that same path eventually leading to leaf node [C], where option 4 is considered due to the status of option 3.

However, subsequent options cannot be considered indefinitely, since the driving time to the request increases with each option. Naturally, despite the dispatch agents being risk averse and preferring subsequent options if there is a potential delay for the closest option, the selected option should still be able to arrive on-time. Since the driving time increases with each option, the driving time, or expected response time, of the furthest option we consider, option 5, is a good indication of whether previous options are able to arrive on-time. This is why multiple nodes testing for the closeness of option 5 to the incident are present in the decision tree, see nodes [D] and [E]. It can be seen that if the closeness of option 5 is sufficiently small, generally lower ranked options are selected for dispatch than when this is not the case.

This is also why the learned model performs significantly better than the limited closest-idle policy in terms of its weighted mean error. In case of sufficient available capacity, dispatch agents clearly prefer risk averse dispatch options. However, while the learned model recognizes that in case of scarcity the dispatch agent is required to choose an ambulance to be dispatched among risky options, the limited closest-idle policy keeps considering subsequent options until a risk-free (completely idle and own region) option is found. In other words, while the performance of the fitted model is similar to the limited closest-idle policy for the majority of dispatch decisions to be made, i.e. in case of sufficient capacity, it strongly outperforms this commonly assumed policy in case of scarce capacity. This ability of the fitted model is especially relevant since dispatch decisions made under scarce capacity are precisely where the expertise and human judgment of the dispatch agents can make a difference.

Fig. 3. Fitted penalty values (on train data) and performance (on test data) of PBCI

4.2 The penalty-based closest-idle policy

The fitted dispatch policy is quite complex. Combined with the fact that a simple model such as the limited closest-idle policy is able to predict dispatch decisions quite well in case of sufficient ambulance capacity, but performs very bad in case of limited capacity due to its inability to consider risky options, leads us to propose a concise, penalty-based policy to represent the dispatch decisions made by BZO's dispatch agents. In line with the three main reasons to deviate from dispatching an ambulance that were deducted from the learned decision tree, penalty terms are defined based on an ambulance's status, region and time until the end of its shift to reflect the potential delay or risk associated with the value of these features. For each ambulance option, its total time penalty is determined based on its status, region and remaining shift time, after which it is added to its driving time. Then, the dispatch option with the lowest driving time plus total penalty is dispatched. In other words, this policy can be called the *penalty-based closest-idle (PBCI) policy*. This approach reflects dispatching agents' preference for a completely idle ambulance from the own region, but ensures that in case of scarce capacity still one of the risky options is selected for dispatch.

These penalty terms are fitted on the training data, such that they result in a maximum weighted F1-score. This is done through an exhaustive search of integer penalty values. The performance of the resulting penalty model is evaluated on the test data. Figure 3 shows the fitted penalty values, the confusion matrix and performance measures. It is shown that both the weighted F1-score and the weighted mean error have improved even further compared to the fitted decision tree. Algorithm 1 shows the resulting PBCI dispatch policy.

The PBCI policy has been presented to and validated by BZO's dispatch agents. Not only did they confirm that the PBCI policy makes sense and is likely to resemble the majority of their dispatch decisions, it also started a constructive discussion on how to improve upon their current decisions. In conclusion, insights from our learned dispatch decision prediction model were used to enrich the commonly assumed closest-idle dispatch policy using penalty values reflecting the

Algorithm 1 Algorithm of the PBCI dispatching policy

```

1: for each dispatch option in the dispatch proposal  $i$  do
2:    $\text{penalty}_i = 0$ 
3:   if ambulance is transferring a patient at a hospital then
4:      $\text{penalty}_i = \text{penalty}_i + 7$  (min.)
5:   else if ambulance is on its way to a less urgent request then
6:      $\text{penalty}_i = \text{penalty}_i + 4$  (min.)
7:   if ambulance is of BNO region then
8:      $\text{penalty}_i = \text{penalty}_i + 6$  (min.)
9:   else if ambulance is of neither BZO nor BNO region then
10:     $\text{penalty}_i = \text{penalty}_i + 10$  (min.)
11:   if shift of ambulance ends within 40 minutes then
12:      $\text{penalty}_i = \text{penalty}_i + 4$  (min.)
13:   Penalized driving time of  $i = \text{driving time of ambulance } i + \text{penalty}_i$ 
14: Dispatch ambulance with smallest penalized driving time

```

risk associated with certain ambulance characteristics. The result of this post-processing phase is a concise policy that has significantly greater resemblance to the actual dispatch decisions made by BZO's dispatch agents compared to the policy that is generally assumed in literature.

5 Current policy as a basis for improvement

The captured dispatch policy provides insight into current practices and gives rise to opportunity for improvement. The PBCI policy provides a basis to improve upon, as well as a benchmark that is close to current practices. To illustrate a possible application of the PBCI policy, we have defined two potential enhancements to current practices and evaluated their potential using a realistic simulation. These enhancements were defined to complement, rather than replace, the current dispatch decision process. Using the PBCI policy as a basis for improvements ensures that practically relevant considerations are included in the improved decision process, fostering adoption. The two potential enhancements to the current dispatch process we propose are (1) *consistently redispersing* ambulances that are on their way to a less urgent request to a more urgent request if this leads to a response time improvement and (2) *reevaluation of active dispatch decisions* upon service completion of an ambulance.

Consistent redispersing From the captured current dispatching process, it can be seen that a dispatch option that is not completely free (on the road or at a station) is considered to be risky due to a potential delay. While a potential delay is difficult to avoid if the ambulance is busy transferring a patient at a hospital, it might be avoided in case of redispersing an ambulance that is currently on its way to a less urgent request. The consistent redispersing policy always dispatch an ambulance that is currently on its way to a less or non-urgent request if this is the best dispatch option for a highly urgent (A1) request. The enhancement is similar to 'reroute-enabled dispatching' as proposed by [8], who evaluated this policy for a hypothetical EMS region consisting of a 16x16 grid, deterministic environment. It is interesting to evaluate the potential performance improvement of

consistently redispatching an ambulance whenever it is the best dispatch option, since the performance improvement might outweigh the disadvantages.

Reevaluation of dispatch decision Currently dispatch decisions are only made upon arrival of a new request. A dispatch decision is made by selecting the best option from those ambulances that are available at that moment. However, the system of ambulances is very dynamic and during the time the dispatched ambulance is driving towards the request, another ambulance may complete serving another request. This other ambulance may in fact be a better dispatch decision than the ambulance that is already on its way. Reevaluation of the dispatch decision might contribute towards improving performance. Contrary to the ‘Parallelism’ dispatch policy of [7], the consideration of a busy ambulance only after it has completed service, prevents dependency on the realization of highly variable treatment times. Furthermore, to prevent reevaluated dispatch decisions resulting in only a marginal difference in response time, as is the case for the ‘free ambulance exploitation’ policy of [8], in our case a reevaluated dispatch decision will only lead to the recently freed ambulance being dispatched instead of the current one if this leads to a response time improvement of at least one minute for highly urgent (A1) requests, or a direct improvement of the on-time performance for less urgent (A2) requests.

5.1 Evaluating potential enhancements using simulation

These two potential enhancements to the current dispatch policy have been evaluated using a realistic simulation that accurately captures the complex dynamics of a real-life size ambulance system within a reasonable computation time. We developed a discrete-event simulation in which the BZO region is aggregated into 138 subregions, corresponding to 4-digit postal codes. Locations of ambulance stations, hospitals, and requests are mapped onto the centroid of its postal code. While our focus is on the performance of urgent (i.e. A1 and A2) requests, we also simulate non-urgent patient transports to capture all dynamics in the utilization of the available ambulance capacity. Furthermore, driving times between each pair of postal codes are assumed to be deterministic, but dynamic, as supplied by the *driving time model* of the RIVM (non-public). The simulation is able to accurately deal with the dynamic arrival of ambulance requests of multiple urgency levels, dynamic ambulance capacity, realistic relocation decisions and a wide range of practical considerations. Furthermore, the captured current dispatch process allowed us to be the first to evaluate alternative dispatch policies by comparing the simulated performance to that of a practically relevant benchmark. The interaction with neighbouring EMS regions was excluded from the simulation due to its complexity. Its effect on the extent to which the simulation resembles reality is expected to be limited due to the fact that external ambulances are rarely dispatched (< 2% of requests).

Table 2 shows that simulating current practices, represented by the PBCI policy, results in a slightly better performance for highly urgent A1 requests and similar performance for moderately urgent A2 requests compared to realized

Table 2. Realized and simulated performance under current dispatch policy

	A1 requests		A2 requests	
	On-time (%)	Mean RT (min:sec)	On-time (%)	Mean RT (min:sec)
Realized	92.13	9:33	97.10	14:32
Simulated	93.63	9:02	97.07	13:38

Table 3. Resulting performance for potential dispatch enhancements

(1) Redispatching (2) Reevaluating	A1 requests		A2 requests		Redispatches/yr. Reevaluations/yr.
	On-time (%)	Mean RT (min:sec)	On-time (%)	Mean RT (min:sec)	
<i>Base</i>	93.63 (+/-0.05)	9:02 (+/-0:00)	97.07 (+/-0.05)	13:38 (+/-0:01)	1425
x	94.06 (+/-0.05)	8:55 (+/-0:00)	96.15 (+/-0.05)	14:04 (+/-0:01)	3293
x	94.04 (+/-0.04)	8:56 (+/-0:01)	97.50 (+/-0.05)	13:34 (+/-0:01)	1413 823
x x	94.40 (+/-0.05)	8:50 (+/-0:00)	96.73 (+/-0.05)	13:58 (+/-0:02)	3269 772

values in the practice of the BZO region. The simulation slightly outperforms reality because the simulation decisions are made consistently, while in practice variations in dispatch decisions occur due to human judgment and differences between dispatch agents. Because we have only a small difference between the realized and simulated performance, we can conclude that our simulation model, with the use of the PBCI policy, is representative for the BZO region.

5.2 Performance of the improved policy

Table 3 shows the resulting performance measures for the two potential enhancements. Besides the main performance measures relating to the response time of urgent requests, the last two columns provide further insight into the effect of both enhancements from which conclusions regarding the effect on ambulance crew disturbance can be deducted. From the effects on performance caused by each dispatch enhancement individually, it can be concluded that both *consistent redispatching* and *reevaluation* of active dispatch decisions upon service completion of an ambulance lead to a significant improvement of the fraction of A1 requests that is served on-time, namely 0.43 and 0.41 percent points (pp) respectively. However, while *consistent redispatching* is quite detrimental for the on-time performance of A2 requests, the *reevaluation* enhancement even improved this measure with 0.43 pp. This detrimental effect of the *consistent redispatching* enhancement on the A2 on-time performance is mostly caused by the fact that an ambulance is redispatched regardless of whether an alternative ambulance is available for dispatch to the original request, and whether this

ambulance is able to arrive on-time. While under the current dispatch policy on average 3.9 redispatches are initiated each day, this number increases to a little over 9 redispatches per day in case of consistent redispaching. Given the number of shifts on an average day, this implies that an ambulance crew is only redispached once every four shifts, which does not seem excessive. From the number of *reevaluations* leading to the recently freed ambulance being dispatched, and thus for the currently dispatched ambulance to be redirected, it can be deducted that such a decision is made on average 2.3 times per day. The disturbance to the ambulance crew of this number of redirections is likely to be quite limited.

We also simulated the combination of enhancements. Adding both the *consistent redispach* and *reevaluation* enhancement yields an even larger performance gain, improving the A1 on-time performance by 0.77 pp. The performance gain of both of these enhancements individually is quite complementary, as combining these enhancements leads to an A1 on-time performance gain of almost the sum of the individuals performance gains. Further, the fact that the *reevaluation* enhancement is beneficial to the performance of A2 requests mitigates part of the detrimental effect of the *consistent redispach* enhancement, leading only to a reduction of 0.33 pp. Combining these two enhancements, however, also leads to a larger number of redirections (resulting from either being redispached or from a reevaluated dispatch decision), which may cause disturbance to ambulance crews. Yet with an average of approximately eleven redirections per day, or an ambulance being redirected once every three eight-hour shifts, this disturbance is likely to be outweighed by the resulting performance gain.

To place the performance gain into perspective, we added additional weekly shifts to the shift roster. The performance gain from our approach is equivalent to adding more than seven weekly ambulance shifts, while our approach does not require additional available resources (see [12] for more details).

6 Conclusion

We captured the way in which dispatch agents currently make decisions on which ambulance to dispatch to a request. Insights from the fitted decision tree were used to enrich the commonly assumed closest-idle dispatch policy, using penalty values reflecting the risk associated with ambulance characteristics. Subsequently, we illustrated an application of the captured dispatch policy by defining two enhancements to current practices and evaluating their performance in a simulation. The proposed approach can be applied to other EMS regions to improve ambulance dispatching.

Dispatch agents in the EMS region Brabant-Zuidoost have indicated to be very happy about the potential of these enhancements to the current dispatch policy in their attempt to push the on-time performance of highly urgent A1 requests to 95%. These process adaptations are essentially free and instantaneous measures to improve performance without increasing the available ambulance capacity. As future research, it is interesting to conduct a field experiment to confirm the potential of the proposed enhancements in practice.

7 Acknowledgements

We would like to thank GGD Brabant-Zuidoost for providing us dispatch data and the overall collaboration. We would like to thank Marko Boon for his help with the development of our simulation model.

References

1. Aringhieri, R., Bruni, M.E., Khodaparasti, S., Van Essen, J.: Emergency medical services and beyond: Addressing new challenges through a wide literature review. *Computers & Operations Research* **78**, 349–368 (2017)
2. Bélanger, V., Ruiz, A., Soriano, P.: Recent optimization models and trends in location, relocation, and dispatching of emergency medical vehicles. *European Journal of Operational Research* **272**(1), 1–23 (2019)
3. Donnot, B., Guyon, I., Schoenauer, M., Panciatici, P., Marot, A.: Introducing machine learning for power system operation support. *arXiv preprint arXiv:1709.09527* (2017)
4. Jagtenberg, C., Bhulai, S., van der Mei, R.: Dynamic ambulance dispatching: is the closest-idle policy always optimal? *Health Care Management Science* **20**(4), 517–531 (2017)
5. Lafond, D., Roberge-Vallières, B., Vachon, F., Tremblay, S.: Judgment analysis in a dynamic multitask environment: Capturing nonlinear policies using decision trees. *Journal of Cognitive Engineering and Decision Making* **11**(2), 122–135 (2017)
6. Lafond, D., Tremblay, S., Banbury, S.: Cognitive shadow: A policy capturing tool to support naturalistic decision making. In: 2013 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA). pp. 139–142. IEEE (2013)
7. Lee, S.: Role of parallelism in ambulance dispatching. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* **44**(8), 1113–1122 (2014)
8. Lim, C.S., Mamat, R., Braunl, T.: Impact of ambulance dispatch policies on performance of emergency medical services. *IEEE Transactions on Intelligent Transportation Systems* **12**(2), 624–632 (2011)
9. Maghrebi, M., Sammut, C., Waller, S.T.: Feasibility study of automatically performing the concrete delivery dispatching through machine learning techniques. *Engineering, Construction and Architectural Management* **22**(5), 573–590 (2015)
10. Maxwell, M.S., Restrepo, M., Henderson, S.G., Topaloglu, H.: Approximate dynamic programming for ambulance redeployment. *INFORMS Journal on Computing* **22**(2), 266–281 (2010)
11. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
12. Theeuwes, N.: Formalization and improvement of ambulance dispatching in Brabant-Zuidoost. Master’s thesis, Eindhoven University of Technology (2019)
13. Van Barneveld, T.: The minimum expected penalty relocation problem for the computation of compliance tables for ambulance vehicles. *INFORMS Journal on Computing* **28**(2), 370–384 (2016)